arXiv:1511.04741v1 [cs.GT] 15 Nov 2015

On the Computational Complexity of Optimal Simple Mechanisms

Aviad Rubinstein*

November 17, 2015

Abstract

We consider a monopolist seller facing a single buyer with additive valuations over n heterogeneous, independent items. It is known that in this important setting optimal mechanisms may require randomization [11], use menus of infinite size [8], and may be computationally intractable [7]. This has sparked recent interest in finding simple mechanisms that obtain reasonable approximations to the optimal revenue [9, 14, 2]. In this work we attempt to find the *optimal simple mechanism*.

There are many ways to define simple mechanisms. Here we restrict our search to *partition mechanisms*, where the seller partitions the items into disjoint bundles and posts a price for each bundle; the buyer is allowed to buy any number of bundles.

We give a PTAS for the problem of finding a revenue-maximizing partition mechanism, and prove that the problem is strongly NP-hard. En route, we prove structural properties of near-optimal partition mechanisms which may be of independent interest: for example, there always exists a near-optimal partition mechanism that uses only a constant number of non-trivial bundles (i.e. bundles with more than one item).

*UC Berkeley. I am grateful to Alon Eden, Amos Fiat, and Muli Safra for inspiring discussions. I also thank Jason Hartline, Christos Papadimitriou, Paul Tylkin, and anonymous reviewers for helpful comments on previous versions of this manuscript. Most of this research was done while the author was an intern at Microsoft Research New England. Part of this research was supported by NSF grant CCF1408635 and Templeton Foundation grant 3966.

1 Introduction

Designing revenue-maximizing mechanisms for a seller who faces a single buyer with additive valuations over n heterogeneous, independent items is a fundamental problem in auction theory. It is known that even in this simple setting, the optimum mechanism requires randomization [11], uses a menu of infinite size [8], and may be computationally intractable [7]. Such mechanisms are often considered "impractical": buyers and sellers may be reluctant to participate in mechanisms that are too complicated; randomization may be restricted by legal requirements (and by our poor understanding of risk aversion); describing and choosing among infinite menus raises obvious issues of communication and computational complexity, etc. Put in computer science jargon, **simplicity is a constraint**: just like the auctioneer cannot obtain the entire social welfare because of incentive compatibility constraints, the optimum mechanism's revenue is infeasible because of the simplicity constraint.

In recent years there have been many works comparing simple mechanisms to optimal mechanisms (including [12, 1, 17, 16, 3]). In particular, a line of works [9, 14, 2] in our setting (a single buyer with additive valuations over independent items) culminated with a celebrated 1/6-approximation of the optimal revenue by the better of the following two mechanisms: (a) sell each item separately; and (b) auction all the items together as one grand bundle. Here, rather than comparing to the benchmark of the globally optimal (but infeasible) auction, we want to find the best feasible mechanism. Clearly, the above mechanism also obtains a 1/6-approximation of the optimal simple mechanism; but can we do better?

Can we find the optimal simple mechanism?

Alas, it is not clear how to formalize "simple". In this work we propose *partition mechanisms* as a standard for simplicity. (In a partition mechanisms the seller partitions the set of items into disjoint bundles, and posts a price for each bundle; the buyer is allowed to select any number of bundles.) In Section 2 we discuss some of the reasons that made us choose this definition, as well as some of its imperfections. We want to emphasize that the same question could be asked with respect to any definition of "simple". (For example: what is the computational complexity of finding the optimal deterministic mechanism with polynomial (additive)-menu-size?)

Our technical contributions include a PTAS, i.e. for any constant $\delta > 0$, we give a polynomial time algorithm that finds a partition mechanism that obtains $(1 - \delta)$-approximation to the optimal revenue among all partition mechanisms. Rather than developing novel algorithmic techniques, our main tool is exploring the structural properties of near-optimal partitions. For example, we prove that there exists a near-optimal partition mechanism with only a constant number of non-trivial bundles. We also prove that this problem is strongly NP-hard, i.e. there is no FPTAS (assuming $\mathsf{P} \neq \mathsf{NP}$).

Organization In Section 2 we discuss some of the merits of partition mechanisms. In Section 4 we give a few interesting examples and provide some intuition for the technical part. The NP-hardness result appears in Section 5 and the PTAS in Section 6.

1.1 Related work

We briefly discuss a few related works on computational complexity of designing simple, revenue-(near)-optimal mechanisms in settings with independent item valuations. For a constant number (or many i.i.d.) additive buyers with monotone hazard rate (MHR) valuation

distributions, Cai and Huang [5] give a PTAS to the optimal mechanism. (Note that we make no assumption on the distributions except independence.) Cai and Huang's mechanism is simple in the sense that most items are sold as a single bundle, but for the remaining few items an arbitrary (potentially randomized) mechanism is used. Our restriction to partition mechanisms for an additive buyer has an analog of item-pricing mechanisms for a unit-demand buyer. For the latter, Cai and Daskalakis [4] give a PTAS for monotone hazard rate valuations and a Quasi-PTAS for regular valuations, and finding the exact optimum for general valuations is NP-hard by Chen et al. [6].

2 Partition mechanisms as simple mechanisms

While there have been notable attempts to quantify complexity of different mechanisms (e.g. by Hart and Nisan [10] and recently by Morgenstern and Roughgarden [15]), it is fair to say that we have not seen an indisputable, universal definition of simple mechanisms. Most likely, because "simple" *can and should* mean different things in different settings; for example, compare the simplicity desiderata in the following scenarios: selling produce in a grocery store (buyers are limited in time and computational capacity); spectrum auctions (buyers may be limited by legal constraints); and ad-auctions in an online marketplace (decisions are often made by automated algorithms).

In this work we define simple mechanisms as partition mechanisms. Certainly, there are issues with this definition. One immediate problem with restricting to partition mechanisms is that they don't really capture all simple mechanisms. In particular, see Example 3 for a distribution where a simple, deterministic mechanism that is not a partition mechanism obtains a strictly greater revenue. More importantly, some of the advantages of partition mechanisms listed in this section are restricted to a single buyer, with additive valuations, over independent items; this issue is illustrated in Example 4 which shows that for many buyers with additive valuations over independent items, partition mechanisms achieve a revenue much lower than the optimum. Coming up with a canonical definition for simple mechanisms remains one of the most important open problems in this line of work. Nevertheless, in this section we argue that partition mechanisms have many advantages as the standard for simple mechanisms in this particular setting.

Expressiveness We argue that despite their simplicity, partition mechanisms can be used to express important auctions of interest. For example, they generalize both selling items separately and bundling all the items together; thus by [2] they guarantee at least a 1/6-approximation to the optimal revenue achievable with any mechanism. Furthermore, this is a strong generalization: as we show in Example 1, partition mechanisms can obtain as much as double the revenue obtained by the better of selling items separately or bundling all the items together. Also, we note partition mechanisms can exhibit rich structure, as is evident by our NP-hardness result.

Menu complexity and false-name-proofness Hart and Nisan [10] discuss a measure of menu-size complexity: every truthful mechanism can be represented as a menu of (potentially randomized) outcomes and prices, where the buyer is allowed to choose one of those outcomes. As noted by Hart and Nisan, the mechanism which auctions each item separately has exponential menu-size complexity under this definition. To overcome this problem, they also introduce a measure *additive-menu-size*, where the buyer is allowed to buy an arbitrary

number of outcomes from the menu. Under this definition, partition auctions have linear additive-menu-size complexity.

A related issue is that of false-name-proofness, i.e. can a buyer gain from participating in the mechanism several times? Partition mechanisms (and additive-menu mechanisms in general) have the advantage that they are always false-name-proof.

Locality and buyer-side computational complexity Partition mechanisms also have the advantage that the buyer's decisions are "local", i.e. the decision to buy one bundle is independent of the decision to buy other bundles. This greatly simplifies tasks such as analyzing and reasoning about such mechanisms, learning or predicting the effects of changes to the environment or the mechanism, etc. In particular, this makes the buyer's decisions very easy.

Revenue monotonicity Hart and Reny [11] observed an interesting phenomenon they call *revenue non-monotonicity*, where increasing the buyer's valuations (in the sense of stochastic dominance), may strictly decrease the optimal obtainable revenue. Hart and Reny showed a constant factor gap between the revenue obtainable with the higher and lower valuations, even when selling two i.i.d. items. Furthermore, [16] recently observed that for two items with correlated valuations, this gap may be infinite. Another nice property of partition mechanisms is that the maximum revenue obtainable by auctions in this class is revenue-monotone.

3 Preliminaries

For any distribution $\vec{D}$ of valuations, we use the following notation, mostly due to [9, 2], to denote the optimum revenue for each class of mechanisms:

- $\textsc{Rev}\left(\vec{D}\right)$ - the maximum revenue among all truthful mechanisms;
- $\textsc{DRev}\left(\vec{D}\right)$ - the maximum revenue among all truthful *deterministic* mechanisms;
- $\textsc{PRev}\left(\vec{D}\right)$ - the maximum revenue among all truthful *partition* mechanisms;
- $\textsc{BRev}\left(\vec{D}\right)$ - the maximum revenue obtainable by auctioning the grand bundle; and
- $\textsc{SRev}\left(\vec{D}\right)$ - the maximum revenue obtainable by pricing each item separately.

When $\vec{D}$ is clear from the context, we simply write $\textsc{Rev}, \textsc{DRev}$, etc.

4 Techniques, intuition, and examples

We begin our technical exposition with the following example which separates the revenue obtainable with a partition mechanism from the better of pricing each item separately or auctioning the grand bundle.

Example 1 ($\textsc{PRev} = (2 - o(1)) \max\{\textsc{SRev}, \textsc{BRev}\}$)**.** Consider $2n$ items:

- A: n items with equal-revenue valuations. $v_a \in S \triangleq \{1, \ldots, \sqrt{n}\}$, with distribution $\Pr[v_a \geqslant k] = 1/k \ \ \forall k \in S$; and

- B: n items with rare-event valuations. $v_b \in \left\{0, \alpha n^b\right\}$, with distribution $\Pr\left[v_b = \alpha n^b\right] = n^{-b}$, where we set $\alpha = \mathsf{E}\left[v_a\right]$.

With a partition mechanism, we can obtain expected revenue $(1 - o(1))\, n\alpha$ from the items in A by bundling them together, and also $n\alpha$ from B by selling each item i separately for price αn^i. However, selling all the items separately achieves negligible revenue on A, whereas the items in B will have negligible contribution to the revenue from selling the grand bundle.

Remark 1. We remark that the Example 1 shows, in particular, a $(2 - o(1))$-gap between $\max\{\textsc{SRev}, \textsc{BRev}\}$ and $\textsc{Rev}$. Previously Babaioff et al. [2] cited an example due to [7] that gave a 1.05-gap.

The example above builds on the key intuition from [14, 2] that there is an interesting tradeoff between bundling and selling separately: when most revenue is distributed among many low-impact, high probability events (as in subset A), their sum concentrates and bundling is preferable; when most revenue comes from rare events (as in subset B), we want to sell the items separately. [14, 2] call this the *core-tail* decomposition.

A nice question suggested to us by Amos Fiat is whether this is the "only way" that $\textsc{PRev}$ can beat $\max\{\textsc{SRev}, \textsc{BRev}\}$. In particular, is there always a revenue-maximizing partition mechanism with at most one non-trivial bundle? The following example shows that the answer is no.

Example 2 (Two non-trivial bundles)**.** Consider the following valuations:

- for $i \in \{1, 2\}$, let $\Pr\left[v_i = 1\right] = \Pr\left[v_i = 2\right] = 1/2$;
- for $i = \{3, 4\}$, let $(1/9) \cdot \Pr\left[v_i = 1\right] = \Pr\left[v_i = 10\right] = 1/10$.

The unique optimal partition mechanism offers bundle $\{1, 2\}$ for price 3 and bundle $\{3, 4\}$ for price 11. The revenue obtained is

$$3 \cdot \Pr\left[\sum_{i \in \{1,2\}} v_i \geqslant 3\right] + 11 \cdot \Pr\left[\sum_{i \in \{3,4\}} v_i \geqslant 11\right] = 3 \cdot \frac{3}{4} + 11 \cdot \frac{19}{100} = 4.34.$$

The core-tail intuition from [14, 2] cannot explain the success of the optimal partition in Example 2. For this distribution, the optimal partition exploits the fact that the values of the bundles are slightly more likely to come out 3 and 11, respectively, than other values on the equal-revenue curve. But they are still far from concentration around 3 and 11. Our NP-hardness result constructs gadgets that generalize Example 2 to create instances where the optimal partition exhibits an arbitrarily complex structure.

Our PTAS is more intricate. Let us informally sketch the main idea. In Example 2, something interesting happens at 3, and something interesting happens at 11. In general, many interesting events can happen in different locations on the (positive) real numbers line, but one of the following two always holds:

- The interesting events are far apart on the real line - in this case we don't lose much by ignoring the events that pertain to lower values. In terms of Example 2, we exploit the asymmetry between the bundle $\{1, 2\}$ and the bundle $\{3, 4\}$.
- Most of the action is restricted to a small interval - this is a redundancy we can exploit. For example, because the sum of many independent random variables in the same range should concentrate.

More concretely, we prove (Lemmata 1 and 2) that there exists a near-optimal partition mechanism that uses only a constant number of non-trivial bundles. In some sense this is a bicreteria-approximation variant of Fiat's conjecture that Example 1 is the only reason we would want to use a partition other $\max\{\text{SRev}, \text{BRev}\}$. We then build on the same intuition to construct modified valuations that approximate the original distributions. Finally, we show that the new distributions admit a succinct representation, so we can find a near-optimal partition mechanism by brute-force search.

For completeness, let us conclude this technical exposition with two examples that separate DRev from PRev; they serve to remind us that there are many interesting mechanisms beyond the scope of partition mechanisms considered in this paper. The first example shows a constant separation in our setting of a single additive buyer.

Example 3 (Hart and Nisan [9]; $\text{PRev} = (1 - \Omega(1))\text{DRev}$)**.** Consider two i.i.d. items with valuations sampled uniformly from $\{0, 1, 2\}$. The expected revenue for selling the bundle with both items (for any price) is at most 1; and selling each item separately (for any price) yields total revenue at most 4/3. Going beyond partition mechanisms, we can offer either item for price 2, or the grand bundle for price 3. The revenue obtained from this auction is

$$3 \cdot \Pr\left[\sum_{i\in\{1,2\}} v_i \geqslant 3\right] + 2 \cdot \Pr\left[\{v_1, v_2\} = \{2, 0\}\right] = 3 \cdot \frac{3}{9} + 2 \cdot \frac{2}{9} = 13/9 > 4/3.$$

The second example shows that with many buyers, partition mechanisms cannot achieve any constant fraction of the optimum revenue. See also the recent paper by Yao [17] on constructing different simple mechanisms in this setting.

Example 4 (e.g. [2]; Many buyers: $\text{PRev} = o(1)\text{DRev}$)**.** We consider n items and $m = n^{1/4}$ buyers; we let v_i^j denote buyer j's value for item i. All v_i^j are drawn i.i.d. from the following distribution: with probability $1 - n^{-3/4}$, $v_i^j = 0$; otherwise, v_j^i is drawn from an equal-revenue distribution with support $\left\{1, \ldots, n^{1/10}\right\}$, i.e. $\Pr\left[v_i^j \geqslant k\right] = n^{-3/4}/k$.

For any one buyer, selling item i for price k yields revenue $n^{-3/4}$, which is only an $O\left(1/\log n\right)$-fraction of the expected value; but the total expected value for the grand bundle concentrates, so that revenue can easily be obtained. With m buyers, the auctioneer can guarantee almost the entire social welfare with the following mechanism: approach buyers in any order; for each buyer charge slightly lower than her expected value for the remaining items, and let her choose her favorite $n^{1/4}$ items. With a partition mechanism, on the other hand, we must fix the partition without knowing which items each buyer wants. Thus partition mechanisms can guarantee at most an $O\left(1/\log n\right)$-fraction of the optimum revenue.

5 NP-hardness

Theorem 1. *Given an explicit description of a product distribution of item valuations, computing a revenue maximizing partition is strongly NP-hard.*

Proof. We reduce from 3D-Matching: Given sets X, Y, Z and a set of hyperedges $H \subseteq X \times Y \times Z$, find a maximum 3-dimensional matching, i.e. maximum non-intersecting subset $M \subseteq H$. Karp [13] proved that it is NP-complete to decide whether there exists a perfect matching (i.e. $|M| = |X| = |Y| = |Z|$).

Construction Identify the set of items with the set of vertices $I \triangleq X \cup Y \cup Z$. Identify the set of hyperedges with their indices; for each $h \in H$, let $\pi_h \triangleq |H|^6 + |H|^3 \cdot h$. Let $\pi_{\min} \triangleq \min_{h \in H} \pi_h = \left|H^6\right|$ and $\pi_{\max} \triangleq \max_{h \in H} \pi_h = \left(1 + O\left(1/|H|^2\right)\right)\left|H^6\right|$, For each item i, the distribution of valuations D_i is defined by the π_h's of its hyperedges. Specifically, we let

$$\operatorname{supp}(D_i) = \{1\} \cup \{\pi_h \colon h \ni i\} \qquad \text{and} \qquad \Pr_{v_i \sim D_i}[v_i \geqslant \pi_h] = 1/\pi_h.$$

Observe that selling each item separately, for any price in its support, yields expected revenue of 1 per item.

Completeness If there exists a perfect matching M, we partition according to this matching, and set the price for bundle h at $\pi_h + 2$. For each bundle, we have

$$\Pr\left[\sum_{i \in h} v_i \geqslant \pi_h + 2\right] = 1 - (1 - 1/\pi_h)^3 = 3/\pi_h - 3/\pi_h^2 + 1/\pi_h^3.$$

The expected revenue for each bundle is therefore

$$(\pi_h + 2)\left(3/\pi_h - 3/\pi_h^2 + 1/\pi_h^3\right) = 3 + 3/\pi_h - O\left(1/\pi_h^2\right)$$

Summing over $|M|$ hyperedges (i.e. $|M|$ bundles), we guarantee a total revenue of $OPT \triangleq |M|\left(3 + 3/\pi_{\max} - O\left(1/\pi_{\max}^2\right)\right)$.

Soundness We first claim that there exists an optimum partition where every bundle is contained in a hyperedge. Let B be a bundle sold for some price $\pi \geqslant \pi_{\min}$. Clearly $\pi \leqslant |B|\,\pi_{\max}$, otherwise it never sells. Similarly, we have $\pi \leqslant |B| + \pi_{\max}$, otherwise it sells with probability at most $|B|/\pi_{\min}^2$, yielding revenue $|B|\,\pi/\pi_{\min}^2 \ll 1$.

Let $i \in B$ be such that $\pi \notin \left[\pi_h - |H|^2, \pi_h + |H|^2\right]$ for all $h \ni i$. We compare the revenue from selling B for price π to the revenue from selling $B \setminus \{i\}$ for price $\pi - 1$. If B sells for price π but $B \setminus \{i\}$ does not sell for price $\pi - 1$, then at least one of the following must be true: (1) $v_i \geqslant \pi - |B|$, which by our assumption on i implies $v_i \geqslant \pi + |H|^2$; or (2) $v_i > 1$ and there is some other $j \in B$ such that $v_j > 1$. We bound the probability of the union as follows:

1. $\Pr\left[v_i \geqslant \pi + |H|^2\right] \leqslant \frac{1}{\pi + |H|^2}$; the revenue loss is bounded by $\frac{\pi}{\pi + |H|^2} \leqslant \frac{\pi - |H|^2/2}{\pi}$;
2. $\Pr\left[(v_i > 1) \wedge (\exists j \in B \;\; v_j > 1)\right] < \frac{|B|}{\pi_{\min}^2}$; the revenue loss is bounded by $\frac{\pi |B|}{\pi_{\min}^2} \leqslant \frac{2|B|}{\pi}$.

There is also some revenue loss from the decrease in price: the original bundle sells with probability at most $\frac{|B|}{\pi - |B|} \leqslant \frac{2|B|}{\pi}$; since we decrease the price by 1, $\frac{2|B|}{\pi}$ also bounds the expected loss in revenue. The total expected loss in revenue is therefore at most $\frac{\pi - |H|^2/2}{\pi} + \frac{4|B|}{\pi} < 1$, so selling B as a bundle cannot be optimal.

There is an optimum partition that bundles items according to hyperedges in some partial matching M', and the rest of the items are in bundles of size at most two. The optimal price for a bundle of two items from hyperedge h is $\pi_h + 1$; the probability of selling for this price is $1 - (1 - 1/\pi_h)^2 = 2/\pi_h - \pi_h^2$. Multiplying by $\pi_h + 1$, we get an expected revenue of $2 + 1/\pi_h$. In particular, this is only $(1 + 1/2\pi_h)$ per item, as opposed to $(1 + 1/\pi_h)$ per item with a full hyperedge. □

6 PTAS

Theorem 2. *For any constant $\delta > 0$, there exists a deterministic polynomial time algorithm that, given an explicit description of a product distribution of item valuations, computes a partition and prices that generate a $(1-\delta)$-approximation to the maximum revenue obtainable by partition mechanisms.*

Proof outline In the next two subsections we prove a structural characterization of near optimal auctions: Lemmata 1 and 2 imply that there exists a near-optimal partition mechanism that uses only a constant number of non-trivial bundles (i.e. bundles with more than one item). Furthermore, the prices to these bundles are all within a constant factor, and all these bundles sell for these prices with constant probability.

In Subsection 6.3 we use our insight about the structure of near-optimal auctions to show that for optimizing over this restricted class of partitions, most of the information in the distribution is redundant. In particular we can place every item in one of $O(\log n)$ buckets, where the items within each bucket are indistinguishable for the algorithm. For each bucket, there are constantly-many options to approximately partition the identical items among constantly many bundles (or to be sold separately). We can enumerate over all approximate partitions for all buckets in polynomial time.

See also description of the algorithm in Subsection 6.4.

6.1 Singletons

Given the following lemma, we can assume wlog that every non-trivial bundle sells with probability at least ϵ.

Lemma 1. *For any $\delta > 0$, let $\epsilon \leqslant \delta^3/4$. Let $B \subseteq [n]$ be a bundle of items, and let $\pi_B \in \mathbb{R}_+$ be an optimal price for B. Suppose that the revenue from auctioning B for price π_B is $\epsilon \cdot \pi_B$, i.e. $\Pr\left[\sum_{i\in B} v_i \geqslant \pi_B\right] = \epsilon$. Then the revenue from selling the items in B separately is at least $(1-\delta)\cdot\epsilon\cdot\pi_B$.*

Proof. For the proof of this lemma, we simplify notation by normalizing to $\pi_B = 1$.

Below we prove that most of the revenue comes from the item with the highest value (this may be a different item in each realization). In particular, if the total value of the bundle is at least 1, then it is likely that there is a single item whose value is almost 1,

$$\Pr\left[\max_{i\in B} v_i \geqslant 1-\delta/2 \mid \sum_{i\in B} v_i \geqslant 1\right] \geqslant 1-\delta/2. \tag{1}$$

This means in particular that $\Pr\left[\max_{i\in B} v_i \geqslant 1-\delta/2\right] \geqslant (1-\delta/2)\,\epsilon$, and therefore selling each item separately for price $(1-\delta/2)$ guarantees a $(1-\delta/2)^2 \geqslant (1-\delta)$-fraction of the revenue from selling B as a bundle.

We now prove (1). Since 1 is an optimal price for B, we have that the $\Pr\left[\sum_{i\in B} v_i \geqslant \delta/4\right] \leqslant \epsilon/(\delta/4)$, otherwise $\delta/4$ would have been a better price. What is the probability that there exist a partition $B = S \cup T$ such that $\sum_{i\in S} v_i \geqslant \delta/2$ and $\sum_{i\in T} v_i \geqslant \delta/2$? If we were to fix any partition $B = U \cup V$ before observing the realizations, or to pick one uniformly at random, we would have

$$\Pr\left[\left(\sum_{i\in U} v_i \geqslant \delta/4\right) \wedge \left(\sum_{i\in V} v_i \geqslant \delta/4\right)\right] \leqslant \left(\epsilon/(\delta/4)\right)^2.$$

Now assume that there exist some partition (S,T) as above, and pick (U,V) uniformly at random. With probability at least $1/4$ we have that $\sum_{i\in(S\cap V)} v_i \geqslant \sum_{i\in(S\cap U)} v_i$ and $\sum_{i\in(T\cap V)} v_i \leqslant \sum_{i\in(T\cap U)} v_i$; and the same for the event that $\sum_{i\in(S\cap V)} v_i \geqslant \sum_{i\in(S\cap U)} v_i$ and $\sum_{i\in(T\cap V)} v_i \leqslant \sum_{i\in(T\cap U)} v_i$. Thus with probability at least $1/2$,

$$\min\left\{\sum_{i\in U} v_i, \sum_{i\in V} v_i\right\} \geqslant \min\left\{\sum_{i\in S} v_i, \sum_{i\in T} v_i\right\}/2.$$

Therefore the probability that there exist such (S,T) is at most $2\left(\epsilon/\left(\delta/4\right)\right)^2$.

Observe that whenever $\sum_{i\in B} v_i \geqslant 1$ and $\max_{i\in B} v_i < 1-\delta/2$, there exists a partition (S,T) as above: Let $S = \{\arg\max_{i\in B} v_i\}$ and $T = B\backslash S$; if $\max_{i\in B} v_i \geqslant \delta/2$, we're done. Otherwise, move items from T to S until $\sum_{i\in S} v_i \geqslant \delta/2$; since the last item we moved from S to T had value at most $\max_{i\in B} v_i < \delta/2$, we have $\sum_{i\in T} v_i \geqslant 1-\delta > \delta/2$. Therefore,

$$\Pr\left[\max_{i\in B} v_i < 1-\delta/2 \mid \sum_{i\in B} v_i \geqslant 1\right] \leqslant \frac{\Pr\left[\left(\max_{i\in B} v_i < 1-\delta/2\right) \wedge \left(\sum_{i\in B} v_i \geqslant 1\right)\right]}{\Pr\left[\sum_{i\in B} v_i \geqslant 1\right]} \leqslant \frac{32\epsilon}{\delta^2}.$$

Plugging in $\delta = 4\epsilon^{1/3}$ yields (1). □

6.2 Bundles

Lemma 2. *For any constants $0 < \epsilon \leqslant 1/2$ and $\delta > 0$ we can replace all the bundles that sell with probability at least ϵ with $\ell = \mathsf{poly}\left(1/\epsilon, 1/\delta\right)$ bundles, while maintaining a $(1-2\delta)$-fraction of the expected revenue.*

Proof. In Claim 1 we show that we can recursively combine bundles until in any interval of multiplicative-constant-length $[\epsilon\cdot\pi, \pi]$, there is at most a constant ($k = 8\epsilon^{-4}\delta^{-3}$) number of bundles. Then, in Claim 2 we show that we can ignore all bundles except those in some slightly larger interval $[\eta\cdot\pi, \pi]$ (for $\eta = \delta^4\epsilon^5\left(1-\epsilon\right)$). This is a union of $\log_\epsilon \eta \leqslant \frac{\log \delta^4\epsilon^4(1-\epsilon)}{\log\epsilon}$ smaller intervals $[\epsilon\cdot\pi, \pi]$; together with the sparsity we obtained in Claim 1, this implies that we are left with at most $\ell = k\log_\epsilon \eta$ bundles. □

Claim 1. For any $\epsilon, \delta > 0$, let $k = 8\epsilon^{-4}\delta^{-3}$. Consider only bundles B_i that sell for price π_i with probability at least ϵ. Partition the positive reals into multiplicative intervals $[\epsilon\cdot\pi, \pi]$. Consider k separate bundles $B_1, \ldots, B_k$ with prices $\pi_1, \ldots, \pi_k$ in the same interval $[\epsilon\cdot\pi, \pi]$, and associated probabilities of selling $p_1, \ldots, p_k$. Whenever we encounter such a k-tuple, we combine them into one bundle B' with price $\pi' = \sum_i p_i\cdot\pi_i$ and probability $p' = 1$. Recurse until every interval has at most k bundles. (The number of bundles decreases at each step, so this process is guaranteed to terminate.) Finally, discount all newly formed bundles by a factor of $(1-\delta/2)$.

Then every newly formed bundle sells with probability at least $1-\delta/2$. In particular, this guarantees an $(1-\delta)$ approximation to the original revenue.

Proof. Let $B' = \bigcup B_i$ be any newly formed bundle, where B_i's, π_i's, and p_i's are the original bundles, prices and probabilities (i.e. B' may denote a union of union of bundles). Let v_{B_i} denote the random variable $v_{B_i} \triangleq \sum_{j\in B_i} v_j$; let also $\hat{v}_{B_i} \triangleq \min\left\{v_{B_i}, \pi_i\right\}$. Then we have

$$\mathsf{E}\left[\sum_i \hat{v}_{B_i}\right] \geqslant \sum_i p_i\cdot\pi_i \quad \text{and} \quad \mathrm{Var}\left[\sum_i \hat{v}_{B_i}\right] \leqslant \sum_i \pi_i^2.$$

Applying Chebyshev's inequality,

$$\Pr\left[\sum_i \hat{v}_{B_i} \leqslant (1-\delta/2)\sum_i p_i\cdot\pi_i\right] \leqslant \frac{\sum \pi_i^2}{\left(\frac{\delta}{2}\cdot\sum_i p_i\cdot\pi_i\right)^2}.$$

In the last union that formed B', we combined at least k bundles, each with $p_i\pi_i \in \left[\epsilon^2\pi,\pi\right]$. Therefore,

$$\frac{\sum_{i\in[k]}\pi_i^2}{\left(\frac{\delta}{2}\cdot\sum_{i\in[k]} p_i\cdot\pi_i\right)^2} \leqslant \frac{\sum_{i\in[k]}\pi^2}{\left(\frac{\delta}{2}\cdot\sum_{i\in[k]}\epsilon^2\pi\right)^2} \leqslant \frac{4}{\delta^2 k\epsilon^4}.$$

Plugging in $k = 8\epsilon^{-4}\delta^{-3} \geqslant -\ln(\delta/2)\cdot(4\epsilon^{-4}\delta^{-2})$ guarantees that we sell the grand bundle with probability at least $1-\delta/2$. □

Claim 2. For any $\epsilon,\delta>0$, and let $\eta = \delta^4\epsilon^5(1-\epsilon)$. Let bundles $B_1,\ldots,B_m$, have optimal prices $\pi_1,\ldots,\pi_m$, and denote $\pi^* = \max_{i\in[m]}\pi_i$. Suppose that bundle B_i sells for price π_i with probability $p_i \geqslant \epsilon$, for every $i\in[m]$. Suppose further that in each range $[\epsilon\pi,\pi]$ of prices we have at most $k=\epsilon^{-4}\delta^{-3}$ bundles (this is wlog by the previous claim). Then a $(1-\delta)$-fraction of the revenue can be obtained by selling only the bundles with with price $\pi_i \geqslant \eta\pi^*$.

Proof. k bundles with prices in interval $[\epsilon\pi,\pi]$ can yield at most $k\pi$ revenue. Summing over $\pi\in\left\{\eta\pi^*,\epsilon\cdot\eta\pi^*,\epsilon^2\cdot\eta\pi^*,\ldots\right\}$, we have that all those bundles together yield revenue at most $k\eta\pi^*/(1-\epsilon)$. Plugging in $\eta=\delta^4\epsilon^5(1-\epsilon)$ completes the proof of the claim. □

6.3 Discretization

In this section we consider a sequence of manipulations on the distribution of each item's valuations. At the end of the manipulation, every item will fit in one of $O(\log n)$ buckets, with all the items in each bucket having indistinguishable distributions. The first step is to discretize the valuation distributions:

Definition 1. Let $\vec{D} \triangleq \times D_i$ be a valuation distribution over non-negative reals ($\mathbb{R}_+$). Let $\mathcal{N}_\epsilon \triangleq \{0\}\cup\left\{\ldots,(1+\epsilon)^{-1},1,(1+\epsilon),(1+\epsilon)^2,\ldots\right\}$ be a multiplicative-$(1+\epsilon)$-net over $\mathbb{R}_+$. For each i, we construct the *rounded valuation distribution* $D_i^{(1)}$ as follows: (a) round down every valuation in the support to the nearest smaller (or equal) element in $\mathcal{N}_\epsilon$; then (b) round down every probability of valuation in the new support to the nearest smaller (or equal) element in $\mathcal{N}_\epsilon$. Finally, we let $\vec{D}^{(1)} \triangleq \times D_i^{(1)}$.

The following lemma implies that for sufficiently small $\epsilon>0$, the loss in revenue from rounding the valuations is negligible.

Lemma 3. *For any constant $\epsilon>0$, non-negative product distribution $\vec{D}$, and price π, we have*

$$\Pr_{\vec{v}^{(1)}\sim\vec{D}^{(1)}}\left[\sum v_i^{(1)} \geqslant (1-\delta)\,\pi\right] \geqslant (1-\delta)\Pr_{\vec{v}\sim\vec{D}}\left[\sum v_i \geqslant \pi\right], \tag{2}$$

where $\delta \triangleq 2\epsilon^{1/3}$.

Proof. Rounding the valuations to $\mathcal{N}_\epsilon$ can decrease the sum by a factor of at most $(1+\epsilon)$. Rounding the probabilities is slightly trickier. An equivalent way of formulating the rounded valuation distribution is to sample $\vec{v}\sim\vec{D}$, round down the valuation of each item to $\mathcal{N}_\epsilon$, and

then zero the valuation of each item independently with probability at most ϵ. Inequality (2) now follows from

$$\Pr_{\vec{v}\sim\vec{D}}\left[\sum v_i^{(1)} \geqslant (1-\delta)\,\pi \mid \sum v_i \geqslant \pi\right] \geqslant 1-\delta.$$

In particular, it suffices to show that for every $\vec{v}$ such that $\sum v_i \geqslant \pi$,

$$\Pr_{\vec{v}^{(1)}}\left[\sum v_i^{(1)} \geqslant (1-\delta)\,\pi \mid \vec{v}\right] \geqslant 1-\delta, \tag{3}$$

where the randomness is only over the independent zeroing of each valuation.

Fix any such $\vec{v}$ and let $\pi_{\vec{v}} \triangleq \sum v_i \geqslant \pi$. The expectation of the sum is at least $\mathsf{E}\left[\sum v_i^{(1)}\right] \geqslant (1-\epsilon)^2\,\pi_{\vec{v}}$, and the variance is at most $\epsilon\cdot(1-\epsilon)^2\,\pi_{\vec{v}}^2$. Therefore by Chebyshev's inequality,

$$\Pr\left[\sum v_i^{(1)} \leqslant (1-\delta)\,\pi_{\vec{v}}\right] \leqslant \frac{\epsilon}{(\delta/2)^2}.$$

Plugging in $\delta = 2\epsilon^{1/3}$ completes yields (3). □

Recall that we can assume wlog that all our bundles sell with constant probability for prices in $[\epsilon\pi, \pi]$ (Lemma 2). Thus, for the purpose of (approximately) evaluating an item's contribution to any bundle it suffices to consider only its valuations in $\left[\pi/n^2, \pi\right]$, rounding down larger valuations to π and ignoring smaller. Notice that the new support has size at most logarithmic: $\left|\mathcal{N}_\epsilon \cap \left[\pi/n^2, \pi\right]\right| = O(\log n)$. Similarly, for the purpose of bundling, we can assume wlog that every valuation in the support has probability at least $1/n^2$. Now each value in the support is associated with one of $\left|\mathcal{N}_\epsilon \cap \left[1/n^2, 1\right]\right| = O(\log n)$ potential probabilities.

In order to represent each item we need to know one more number - the revenue it can generate when sold separately. Here again we can assume wlog that this revenue is in $\left[\pi/n^2, \pi\right]$: if it is less than π/n^2, its revenue is negligible and we never want to sell this item separately; if it is greater than π, we always want to sell this item separately. The revenue from selling an item separately is a product of two numbers (price and probability) in $\mathcal{N}_\epsilon$, and therefore also belongs to $\mathcal{N}_\epsilon$. As before, this means that we can assume wlog that the expected revenue takes one of $\left|\mathcal{N}_\epsilon \cap \left[\pi/n^2, \pi\right]\right| = O(\log n)$ values.

So far for each item we need $O(\log n)$ numbers, each from a set of size $O(\log n)$. While this is much more succinct than the naive representation, it is still not good enough for our algorithmic application (at this point we still need $(\log n)^{O(\log n)} \gg n$ buckets). In the next two steps we reduce to only three numbers from sets of size $O(\log n)$: first, we show that in the lower end of the support it suffices to keep the aggregate expectation rather than probability of each value; second, we argue that we can assume wlog that all the high values in the support have approximately the same probability; and the third number is the expected revenue from selling separately.

Low values

Fix any $\epsilon > 0$, distribution $D_i^{(1)}$ over $\mathbb{R}_+$, and $\pi \in \mathbb{R}_+$. Define

$$v_i^{(2)} \triangleq \begin{cases} v_i^{(1)} & v_i^{(1)} \geqslant \epsilon\pi \\ \mathsf{E}_{u\sim D_i^{(1)}}\left[u \mid u < \epsilon\pi\right] & v_i^{(1)} < \epsilon\pi \end{cases};$$

round down the new value and probability to the nearest smaller elements in $\mathcal{N}_\epsilon$, and let $\vec{D}_i^{(2)}$ denote the resulting new distribution.

Intuitively, it may be helpful to think of $\vec{D}^{(2)}$ as "erasing" or "blurring" the information about $\vec{D}^{(1)}$ below $\epsilon\pi$. Notice that for any $\epsilon > 0$, $\mathsf{E}_{v_i^{(2)} \sim D_i^{(2)}}\left[v_i^{(2)}\right] = \mathsf{E}_{v_i^{(1)} \sim D_i^{(1)}}\left[v_i^{(1)}\right]$. The next lemma shows that $\vec{D}^{(2)}$ also generates approximately the same revenue.

Lemma 4. *Let $\epsilon > 0$, let $\delta = 2\epsilon^{1/3}$, and let $\vec{D}^{(1)}$ be a product distribution. Then,*

$$\Pr_{\vec{v}^{(1)} \sim \vec{D}^{(1)}}\left[\sum v_i^{(1)} \geqslant \pi\right] - 2\delta \leqslant \Pr_{\vec{v}^{(2)} \sim \vec{D}^{(2)}}\left[\sum v_i^{(2)} \geqslant (1-\delta)\,\pi\right] - \delta \leqslant \Pr_{\vec{v}^{(1)} \sim \vec{D}^{(1)}}\left[\sum v_i^{(1)} \geqslant (1-2\delta)\,\pi\right].$$

Recall that by Lemma 1, we can assume wlog that all bundles sell with probability $\Omega(1)$; thus we can tolerate the above additive loss in probability.

Proof. By Chebyshev's inequality, the sum of valuations less than $\epsilon\pi$ is within an additive $\pm\delta\pi$ of its expectation with probability at least $1 - \delta$ (and the rest of the valuations don't change at all). □

High values

For any $\epsilon > 0$, a rounded distribution $D_i^{(2)}$ (in the sense of Definition 1), and $\pi \in \mathbb{R}_+$, let $p_i^* \triangleq \max_{v \in [\epsilon\pi, \pi]} \Pr_{u \sim D_i^{(2)}}[u = v]$ denote the most likely valuation in $[\epsilon\pi, \pi]$. Let $L^{\epsilon}_{D_i^{(2)}}$ denote the set of unlikely high valuations:

$$L^{\epsilon}_{D_i^{(2)}} \triangleq \left\{ v \colon \Pr_{u \sim D_i^{(2)}}[u = v] < \epsilon^4 p_i^* \right\} \cap [\epsilon\pi, \pi]$$

Let $D_i^{(3)}$ denote the restriction of $D_i^{(2)}$ to $\operatorname{supp}\left\{D_i^{(2)}\right\} \setminus L^{\epsilon}_{D_i^{(2)}}$:

$$v_i^{(3)} \triangleq \begin{cases} v_i^{(2)} & v_i^{(2)} \notin L^{\epsilon}_{D_i^{(2)}} \\ 0 & v_i^{(2)} \in L^{\epsilon}_{D_i^{(2)}} \end{cases}.$$

The following lemma implies that for each i it suffices to maintain p_i^* (which always takes one of $\left|\mathcal{N}_\epsilon \cap \left[1/n^2, 1\right]\right| = O(\log n)$ values), and a constant number of bits for each of the (constantly many) values in $\operatorname{supp}\left\{D_i^{(3)}\right\} \cap [\epsilon\pi, \pi]$.

Lemma 5. *Let $\vec{D}^{(2)}$ be a rounded product distribution of valuations of items in B. Assume that selling bundle B for price π yields higher expected revenue than selling all the items in B separately. Then,*

$$\Pr_{\vec{v}^{(3)} \sim \vec{D}^{(3)}}\left[\sum_{i \in B} v_i^{(3)} \geqslant \pi\right] \geqslant (1-\delta) \Pr_{\vec{v}^{(2)} \sim \vec{D}^{(2)}}\left[\sum_{i \in B} v_i^{(2)} \geqslant \pi\right].$$

Proof. We compare, for each item i, the potential revenue loss from switching from $D_i^{(2)}$ to $D_i^{(3)}$ to the expected revenue from selling i separately. We argue that for each i, the latter is much larger. Summing over all items, we have that the total loss is much smaller than the total revenue from selling every item separately. By the premise, the latter is less than the original revenue from selling the bundle.

By definition, item i has probability at least p_i^* of having value $\epsilon\pi$. Thus we can obtain revenue at least $p_i^* \cdot \epsilon\pi$ from selling it for price $\epsilon\pi$. In contrast, every time we zero a valuation, we could potentially lose revenue π. The total probability over items in $L^{\epsilon}_{D_i^{(2)}}$ is $\left|L^{\epsilon}_{D_i^{(2)}}\right| \epsilon^4 p_i^*$. Between $\epsilon\pi$ and π, there are at most $-\log\epsilon/\log(1+\epsilon) < 1/\epsilon^2$ elements in $\mathcal{N}_\epsilon$; in particular, $\left|L^{\epsilon}_{D_i^{(2)}}\right| < 1/\epsilon^2$. Therefore, the revenue lost is at most an ϵ-fraction of the revenue from selling i separately. □

From QPTAS to PTAS

We have reduced the representation of each item to three numbers in $\mathcal{N}_\epsilon \cap \left[\pi/n^2, \pi\right]$: the expectation over lower values, $\mathsf{E}_{u\sim D_i^{(3)}}\left[u \mid u < \epsilon\pi\right]$ (rounded down to $\mathcal{N}_\epsilon$); the quantity $p_i^* \cdot \pi$, where $p_i^* \triangleq \max_{v\in[\epsilon\pi,\pi]} \Pr_{u\sim D_i^{(3)}}\left[u = v\right]$ is the maximum probabilities over higher values; and the maximum revenue from selling i separately, $\textsc{SRev}(D_i)$. (For the higher values we also need a constant number of bits to specify which values have probabilities close to p_i^*, and how close.) At this point we need $O\left(\log^3 n\right)$ buckets, which would suffice for obtaining a Quasi-PTAS.

Our final step is to observe that if any of those three numbers is much lower than the maximum of the three, it might as well be zero. If $\mathsf{E}_{u\sim D_i}\left[u \mid u < \epsilon\pi\right]$ is much higher than $p_i^* \cdot \pi$, then it's contribution to revenue for any bundle outweighs the contribution from any of the higher values appearing with very low probability; similarly if $\mathsf{E}_{u\sim D_i}\left[u \mid u < \epsilon\pi\right] \gg \textsc{SRev}(D_i)$ we would always sell item i as part of one of the bundles. If $\textsc{SRev}(D_i^\epsilon)$ is much higher than either of the other two, then the contribution from selling item i separately outweighs the contribution (from the lower values, higher values, or both) to the revenue from any bundle. Finally, since $\textsc{SRev}(D_i) \geqslant p_i^* \epsilon\pi$, the revenue from selling separately is never much lower than the higher values' contribution.

6.4 Algorithm

We achieve a $(1-\delta)$-approximation of the optimal partition revenue for some constant $\delta > 0$; let $\epsilon = \epsilon(\delta) > 0$ be a sufficiently small constant, and let

$$\mathcal{N}_\epsilon \triangleq \{0\} \cup \left\{\ldots, (1+\epsilon)^{-1}, 1, (1+\epsilon), (1+\epsilon)^2, \ldots\right\}.$$

For each item i, compute the optimum expected revenue from selling i separately, $\textsc{SRev}(D_i)$.

Before we analyze bundles, we first want to guess a range $[\epsilon\pi, \pi]$ in which all the bundle prices will lie. Let $v^{\min}$ denote the minimum over all nonzero values in the support of all items, and let $v^{\max}$ denote the sum, over all items, of the maximal values in their supports. An optimal π must belong to $\left[v^{\min}, v^{\max}\right]$. Enumerate over all potential π's in $\mathcal{N}_\epsilon$. By Lemma 2 for some choice of π, it suffices to optimize only over partitions with a constant number of non-trivial bundles, and each of those bundles sells for prices in $[\epsilon\pi, \pi]$ with probability at least ϵ. For the rest of the algorithm assume we have such an optimal choice of π.

For each i, round D_i as in Definition 1, and let $\vec{D}^{(1)}$ denote the resulting product distribution. ($\vec{D}^{(1)}$ is stochastically dominated by $\vec{D}$, thus the revenue obtained from a partition mechanism with valuations drawn from $\vec{D}^{(1)}$ is at most the revenue obtained with the same partition and pricing with valuations drawn from $\vec{D}$. In the other direction, Lemma 3 guarantees that the revenue lost is at most a small constant fraction.)

For each i, replace all values in $\left[0, \epsilon^2\pi\right]$ with their expectation, and let $\vec{D}^{(2)}$ denote the resulting product distribution. (By Lemma 4, optimizing over bundles with prices in $[\epsilon\pi, \pi]$ with valuations drawn from $\vec{D}^{(2)}$ is the same up to a small constant factor as with valuations drawn from $\vec{D}^{(1)}$.)

For each i, let p_i^* denote the maximal probability $D_i^{(2)}$ gives to any value in $\left[\epsilon^2\pi, \pi\right]$. Remove values with probabilities much smaller than p_i^* from the support of $D_i^{(2)}$, and let $\vec{D}^{(3)}$ denote the resulting product distribution. ($\vec{D}^{(3)}$ is stochastically dominated by $\vec{D}^{(2)}$, thus the revenue obtained from a partition mechanism with valuations drawn from $\vec{D}^{(3)}$is at most the revenue obtained with the same partition and pricing with valuations drawn from $\vec{D}^{(2)}$. In the other direction, Lemma 5 guarantees that the revenue lost is at most a small constant fraction.)

For each i, we now have three variables which may be of different scale: $\textsc{SRev}\left(D_i\right)$, $\mathsf{E}_{u\sim D_i^{(3)}}\left[u \mid u < \epsilon^2\pi\right]$, and $p_i^* \cdot \pi$; we also have the full description of $D_i^{(3)}$ restricted to $\left[\epsilon^2\pi, \pi\right]$, which given p_i^* requires only a constant number of bits. If any of the three variables is much smaller than any of the others, set the smaller variable to zero.

We now represent each item with a constant number of variables, which are all either zero or within constant factors. In total, we have at most $O\left(\log n\right)$ distinct representations, which we henceforth call *buckets*.

Enumerate over the number of bundles (by Lemma 2 it suffices to consider only numbers up to some constant ℓ). For each bucket, we must decide how many items to allocate to each of the ℓ bundles, and which to sell separately. I.e. we must pick some vector in $[0,1]^{\ell+1}$, and up to $\pm\epsilon$, there are at most $\epsilon^{-(\ell+1)}$ different vectors; in particular, only a constant number of choices for each bucket. Enumerate (in polynomial time) over all choices for all $O\left(\log n\right)$ buckets.

References

[1] Saeed Alaei, Hu Fu, Nima Haghpanah, and Jason Hartline. The Simple Economics of Approximately Optimal Auctions. In *the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2013.

[2] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. A simple and approximately optimal mechanism for an additive buyer. *FOCS*, 2014.

[3] MohammadHossein Bateni, Sina Dehghani, MohammadTaghi Hajiaghayi, and Saeed Seddighin. Revenue Maximization for Selling Multiple Correlated Items. *arxiv report*, 2015. http://arxiv.org/abs/1412.3187.

[4] Yang Cai and Constantinos Daskalakis. Extreme-Value Theorems for Optimal Multidimensional Pricing. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, pages 522–531, 2011.

[5] Yang Cai and Zhiyi Huang. Simple and Nearly Optimal Multi-Item Auctions. In *the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2013.

[6] Xi Chen, Ilias Diakonikolas, Dimitris Paparas, Xiaorui Sun, and Mihalis Yannakakis. The complexity of optimal multidimensional pricing. In *the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2014.

[7] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. The Complexity of Optimal Mechanism Design. In *the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2014.

[8] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Strong duality for a multiple-good monopolist. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*, pages 449–450, 2015.

[9] Sergiu Hart and Noam Nisan. Approximate revenue maximization with multiple items. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, pages 656–656, 2012.

[10] Sergiu Hart and Noam Nisan. The menu-size complexity of auctions. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, EC '13, pages 565–566, 2013.

[11] Sergiu Hart and Philip J. Reny. Maximal Revenue with Multiple Goods: Nonmonotonicity and Other Observations. Discussion Paper Series dp630, The Center for the Study of Rationality, Hebrew University, Jerusalem, November 2012.

[12] J.D. Hartline and T. Roughgarden. Simple versus optimal mechanisms. In *EC'09*, pages 225–234, 2009.

[13] Richard M. Karp. Reducibility among combinatorial problems. In *Proceedings of a symposium on the Complexity of Computer Computations, held March 20-22, 1972, at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York.*, pages 85–103, 1972.

[14] Xinye Li and Andrew Chi-Chih Yao. On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences*, 110(28):11232–11237, 2013.

[15] Jamie Morgenstern and Tim Roughgarden. The pseudo-dimension of near-optimal auctions. *CoRR*, abs/1506.03684, 2015.

[16] Aviad Rubinstein and S. Matthew Weinberg. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*, pages 377–394, 2015.

[17] Andrew Chi-Chih Yao. An n-to-1 Bidder Reduction for Multi-Item Auctions and its Applications. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 92–109, 2015.